

Design Issues of JPQ: a Pattern-based Query Language for Document Databases

Xuhui Li¹, Mengchi Liu¹, Xiaoying Wu¹ and Jinguang Gu²

¹ Computer School, Wuhan Univ.,
Wuhan, China

lixuhui@whu.edu.cn, charleslzq@gmail.com, mengchi@sklse.org,
xiaoying.wu@gmail.com

² School of Computer Sci. and Tech., Wuhan Univ. of Sci. and Tech.,
Wuhan, China
simon@wust.edu.cn

Abstract. Document databases are becoming popular, but how to present complex document query to obtain useful information from the document remains an important topic to study. In this paper, we describe the design issues of a pattern-based document database query language named JPQ. JPQ uses various expressive patterns to extract and construct document fragments following a JSON-like document data model. It adopts tree-like extraction patterns with a coherent pattern composition mechanism to extract data elements from hierarchically structured documents and maintain the logical relationships among the elements. Based on these relationships, JPQ deploys a deductive mechanism to declaratively specify the data transformation requests and considers also data filtering on hierarchical data structure. We use various examples to show the features of the language and to demonstrate its expressiveness and declarativeness in presenting complex document queries.

1 Introduction

Document databases are a kind of so-called NoSQL databases which use nested collections of records or rows to store information. The data in the documents are organized in an *aggregate*-oriented way where an aggregate is a collection of related objects. An aggregate is treated as a unit of manipulation meanwhile it has allowable structure and types to enable flexible accessing. Document databases often organize large amounts of aggregates by associating them with the keys to be accessed efficiently and semantically. The dominant document format is JSON[1] which is widely used as a lightweight data-interchange language.

A JSON-like document database is quite different from conventional key-value stores where the aggregates are usually opaque blobs of meaningless bits to be parsed and processed by applications. It is also different from XML database, even though they both exhibit a hierarchical style and they both use semantic labels, i.e., tags or keys, to associate contents, because the keys in JSON documents are often unique in an object whereas different sub-elements of an XML

element can be of same tags. Additionally, JSON has important performance and resource utilization advantages over XML[2], which makes the JSON-like document databases particular interesting.

In recent years some document databases like MongoDB[3], CouchDB[4], OrientDB[5] and RavenDB[6] have been developed and are becoming popular in practice. With the wide spread use of document databases, how to query JSON-like documents has become an important issue. Existing document databases usually provide lightweight query interfaces to satisfy basic query needs. For instance, MongoDB uses certain APIs, such as the “*find()*” function with the keys and constraints as the arguments, to extract the corresponding elements from source documents in the runtime. CouchDB develops a view-based query interface where the predefined views are key-value pairs generated by underlying map/reduce functions. OrientDB supports a simple subset of SQL to extract the document fragments in collection of records, besides its native java-based query API. RavenDB presents its queries with LINQ programs on .NET framework, using the classes developed on indices. Although the above databases allows simple data extraction from documents, their query APIs are not expressive enough for complex requests. For example, to extract multiple interdependent data elements or to construct a specified data structure, users often have to resort to host programming languages. Therefore, a full-fledged query language is necessary to facilitate presenting complex document queries.

Some researchers have begun to study query languages for JSON-like documents recently. Jaql[7] is a functional query language that enables users with function pipelines to extract, filter, join and group JSON data. It enables computation on data collections with the *transform* command and allows some simple structure manipulation with the commands like *expand*. JSONiq[8] is a query language which adopts a subset of XQuery syntax and semantics to query JSON documents, by treating JSON documents as simple XML documents. Similarly, some studies [9,10,11] also employ XPath-like expressions to extract JSON documents in a navigational way, but they have no coherent and complete filtering and constructing mechanism and thus cannot be called a full-fledged language. UnQL[12] is a SQL-like JSON query language which enables selecting and manipulating object elements in data hierarchy with navigation operators. These languages are more expressive than the aforementioned query interfaces, but they are not expressive and declarative enough to present complex queries. More specifically, since they often use navigation operators on documents to extract homogeneous data as plain tuples, it is not easy for them to extract multiple interdependent elements with inherent structure, to handle heterogeneous data and to construct hierarchical document structure. These common requests have to be carried out by multiple extraction commands, nested loop statements and other specific mechanisms. It requires the users to consider the detailed procedures of queries themselves and to be skillful in programming.

In this paper, we present our preliminary study on a pattern-based query language named JPQ (standing for JSON-like Pattern Query). JPQ is a functional language which is inspired by our former study on declarative XML queries[13].

It adopts expressive patterns with coherent mechanisms to extract multiple document data elements simultaneously and construct document fragments in a flexible structure, and thus enabling complex document queries be presented in a declarative way. Our work makes the following contributions: a) Although some pattern-based query languages were once studied on XML, to our knowledge our work is the first attempt to use tree-like patterns in querying JSON style document databases. b) In comparison with existing studies on document query, JPQ focuses on the expressiveness of the query language and adopts various kinds of data extraction patterns to facilitate presenting complex document queries which are useful in practice. c) JPQ deploys a deductive rewriting mechanism to specify data transformation requests, which allows complex data construction to be presented declaratively.

The remainder of the paper is organized as follows. In order to present our query language, we first introduce a simple model of JSON-like documents in Section 2, together with a brief discussion on the schema-less feature and its effects on document database queries. Section 3 introduces the JPQ language with certain examples demonstrating its usage and features. Section 4 compares JPQ and other document query languages or interfaces and concludes the paper.

2 A Simplified Model for Document Query

JPQ adopts a simplified model of JSON-like documents named JHM, standing for JSON-like Hierarchy Model. JHM uses the following grammar to specify the data structure of key-value document fragments in JSON style. As the grammar shows, a JHM document fragment is represented as a value which can be an atomic one like a string or a number, an array of values, or an object comprising one or more key-value pairs. A JHM document often exhibits a data hierarchy as the values in objects or arrays are expanded. JHM requires the string-based keys to be unique in an object, as document database usually does.

```

<value> ::= <atom> | <array> | <object>
<atom> ::= <string> | <number> | <boolean> | empty | ...
<array> ::= [ <value> ( ; <value>)* ]
<object> ::= {<keyvalue> ( , <keyvalue>)* }
<keyvalue> ::= <string>:<value>

```

Although JSON format is evolved from the object-oriented paradigm, document data as JHM specifies are different from conventional object data. The fundamental distinction is that JHM adopts no compound data type and, further, no predefined data schema. This schema-less property essentially affects document queries in several ways. Firstly, it disables randomly accessing an arbitrary part of an object, thus a certain enumeration mechanism is required to parse the object and get the details. Secondly, it does not use labels with a predefined semantics, thus querying key strings becomes an important way to find expected information, because the key strings with the self-described semantics play the role of semantic labels of the associating values. Finally, it supports the

heterogeneity of array elements which are often assumed to convey similar information, thus users often need to manually align heterogeneous array elements to produce homogeneous results.

The following is a document fragment *univ* illustrating a scenario of a university human resource information system where the information of the presidents, the schools and the faculties is stored in a hierarchical document.

```
{ "president":{"ID":"0001", "last name":"Li", "first name":"XH",
  "email":"xxli@123.edu"},
  "executive-vice-president":{"ID":"0002","last name":"Feng",
    "firstname":"YM", "email":xxfeng@123.edu"},
  "vice-presidents":[{"ID":"0003","surname":"Zhou","givenname":"CB"};...]
  "schools":[{"name":"Computer School", "dean":{"ID":"0011", ...},
    "faculty":[{"ID":"0001", "first name":"Li",...,
      "email":"xxli@cs.123.edu:"};
      ...}],
    ... ],
  ... }
```

(univ)

We consider the following three queries on the human resource information: a) find all presidents' personal information; b) for each president find the school he/she belongs to; and c) for each faculty member list the schools in which he/she has an occupation. These queries are common in practical information system. However, if a user tries to present the queries using existing document query interfaces or languages, he might have to spend a lot of time to write tedious programs to extract interdependent data elements, to handle heterogeneity and to transform hierarchical data structure. So here comes JPQ.

3 The JPQ Language

3.1 Synopsis of JPQ

JPQ is a declarative query language which deploys hierarchical patterns to specify queries on JHM documents. It stems from our previous work on XML query but is considerably modified to cater for the practical document query requests mentioned above. A common JPQ program is composed of the **from**, the **construct** and the **where** clauses to present data extraction, construction and filtering respectively. The synopsis of the JPQ grammar is listed below.

```
<query> ::= from <doc> <extraction-pattern>
          (,<doc> <extraction-pattern>)*
          construct <construction-pattern>
          where <conditions>
```

A JPQ query works as follows. The *from* clause in a JPQ program uses one or more extraction statements, i.e., an extraction pattern following a source document, to specify extraction requests. The data elements in the document

are extracted if their surrounding elements and contexts match the corresponding parts of the pattern. These data elements are naturally organized following the structure derived from the extraction pattern. This structure is coherently restructured and used in the construction pattern of the *construct* clause to transform the extracted data elements and generate the output results. Additionally, if there is a *where* clause, the extracted data elements should be filtered by the conditions specified in it before the construction of the final results.

3.2 Data Extraction

Extraction patterns in JPQ include key-value patterns and value patterns which are to respectively match the key-value pairs and the values in JHM documents. We use p_k and p_v to denote the key-value pattern and the value pattern, use v to denote the variables, use r_s and r_v to denote the string predicate and the value predicate, and list the abstract syntax of the extraction patterns as below:

$$\begin{aligned} p_k &::= v:p_v \mid r_s:p_v \mid (v \ r_s):p_v \mid *:p_v \mid p_k|p_k \\ p_v &::= v \mid r_v \mid * \mid \{p_k, \dots, p_k\} \mid [p_v] \mid <p_v, p_v> \mid p_v|p_v \mid /p_k \mid //p_v \end{aligned}$$

As the syntax rules show, an extraction pattern is often structurally composed of the variables, the predicates and the wildcard “*” with the structural operators like “:”, “{ }” or “[]” and the logical operators like “<, >” or “|”. The variables and the predicates for the key string or value are to test whether a document fragment, i.e., a key or a value, satisfies the restrictions indicated by the pattern, and if it does, to bind the fragment to the variable (if there is one) in the pattern. For example, to match a key-value pair “‘executive-vice-president’:v” with a key-value pattern “\$k ‘?president?’:*”, the key “executive-vice-president” is tested by the string predicate “?president?” for the existence of a substring “president”, and \$k would be bound to the key, denoted as the matching pair “\$k→‘executive-vice-president’”. The value v is ignored as it is matched with the wildcard “*”. On the other hand, matching the element with the key-value pattern “‘?president?’:\$x” would result in a matching pair “\$x→v”.

For an extraction pattern which can bind various parts of a document fragment to one or more variables, the matching pairs like “\$x→v” are organized structurally according to the composite structure of the pattern. JPQ adopts a logical way to compose the matching pairs by introducing three kinds of structures namely **tuple**, **array** and **option**.

A tuple of the form (r_1, \dots, r_n) conjunctively combines the subordinate matching results r_1, \dots, r_n . It is generated in matching a fragment with a multi-variable pattern, such as a definitive key-value pattern, an object pattern or a conjunctive pattern, where the variables are considered to be conjunctively associated. JPQ also introduces the conjunctive pattern $<p_1, \dots, p_n>$ for the queries where a value is to be respectively matched with the patterns p_1, \dots, p_n and the results are combined as a tuple. Here the pattern p_i can also be a predicate required to be satisfied by the value. For example, matching the aforementioned key-value with the pattern “\$k‘?president?’:<\$p, {‘last name’: <\$l, ‘F?’>}>” would generate a tuple “(\$k→‘executive-vice-president’, \$p→v, \$l→‘Feng’)”.

An array of the form $[r_1; \dots; r_n]$ is an ordered list of the results r_1, \dots, r_n , which derives from matching fragments with an array pattern or an enumeration pattern. For example, matching the “*vice-presidents*: $[v_1; \dots; v_n]$ ” in *univ* (where v_1, \dots, v_n denotes the values in the array) with the pattern “*?president?*: $[\$p]$ ” will result in the array “ $[\$p \rightarrow v_1; \dots; \$p \rightarrow v_n]$ ”.

JPQ introduces the enumeration patterns for generating an array of matching results from an arbitrary fragment, borrowing the convenient path operators “/” and “//” from XML query languages. These patterns are used to parse and enumerate the content of a fragment when its inner structure is unknown, as previously mentioned. The children pattern $/p$ is to match an object value by iterating its elements and matching them with the pattern p . The matching results of p would form an array following the document order. For example, matching *univ* with the pattern “ $/\$r \text{ '?president?':*}$ ” will result in the array “ $[\$r \rightarrow \text{'president'}; \$r \rightarrow \text{'executive-vice-president'}; \$r \rightarrow \text{'vice-presidents'}]$ ”. The descendants pattern $//p$ is to recursively find all the values matching the pattern p under current context and the sub-contexts. The matching results would form an array following the preorder-traversal order.

An option derives from matching a fragment with an optional pattern $p_1 | p_2 \dots | p_n$ which provides multiple choices of pattern matching. The fragment would be disjunctively matched with the patterns p_1, \dots, p_n , and the results would be filtered by the conditions in the *where* clause if such conditions exists. After that, the first valid result of p_i would be used as the result of the whole pattern. Option patterns are useful in handling heterogeneity. For example, matching a person value of *univ* with the pattern “*first name*: $\$f_1 | \text{'firstname':}\$f_2 | \text{'given name':}\g ” will return his/her first name or given name, and thus heterogeneous person names can be unified and processed homogeneously; on the other hand, matching a name value with the pattern “ $\text{'<\$n1, 'L?>' | <\$n2, 'W?>'}$ ” will bind the value to different variables depending its initial letter, and thus homogeneous data can be separated and handled heterogeneously.

Using the extraction patterns described above, JPQ can flexibly and expressively extract and organize the elements. For example, to find the schools where a president works as a faculty member, we can use the following extraction pattern to find the related information.

```
doc("univ") </\$r'?president?':(<\$p1,{ "ID":\$id1}>|<\$p2,{ "ID":\$id2}>)>,
    {"schools":[{"name":\$n, "faculty":[{"ID":\$id3}]}]}>
```

This extraction pattern is a conjunctive one composed of the enumeration pattern “ $/\$r \text{ '?president?':*}$ ” and the object pattern “ $\{\text{'schools':...}\}$ ”. In matching the document with the former pattern, all the key-value pairs in the object would be enumerated and matched with the key-value pattern. The key-value pairs whose key contains the substring “*president*” would be bound to $\$r$ and its value, denoted as v_1, v_2, \dots , would be matched with the option pattern “ $\text{'<\$p1...>' | <\$p2...>'}$ ”. In matching the document with the latter pattern, the array value of the “*schools*” would be matched with the succeeding array pattern to extract each school’s name and IDs of its faculty members. Part of the matching result is listed as follows:

```
( [ (\$r -> "president", \$p1 -> v1, \$id1 -> "0001");
```

```

($r -> "executive-vice-president", $p1 -> v2, $id1 -> "0002");
($r -> "vice-presidents", [($p2 -> v3, $id2 -> "0003"); ... ]),
[ ($n -> "Computer School", [$id3 -> "0001"; ...]); ... ] )

```

3.3 Data Construction

Data construction in JPQ is presented with construction patterns which are essentially the function invocations on the **matching terms**. A matching term is an expression to specify the (transformed) structure of the matching results of an extraction pattern. The abstract syntax of matching term is

$$t ::= v \mid (t, t) \mid t|t \mid [t]_t \mid \hat{[t]}_t \mid t\%$$

As the syntax shows, a matching term is often composed of the variables in the result and the operators like “(,)”, “[]”, “|” and “%”, and thus is called a tuple, array, option or distinct term. For the array term $[t]_{t'}$, t' is named the **index term** of t and t can usually be specified as a (t', t) . It is required that in the array of matching results with the term $[t]_{t'}$, each item of t can be identified by an item of t' . Generally, the element of an array of the original matching results should be indexed by itself, that is, its matching term should be like $[t]_t$. For the distinct term $t\%$, it denotes a distinct value in a given array, which is introduced in detail later. For example, a tuple result “ $(\$x \rightarrow v_1, \$y \rightarrow v_2)$ ” is of the tuple term “ $(\$x, \$y)$ ”, and an array “ $[\$x \rightarrow v_1; \$y \rightarrow v_2; \$x \rightarrow v_3]$ ” is of the array term “ $[\$x|\$y]_{\$x|\$y}$ ”. The matching term of an extraction pattern p is often denoted as a function “ $mt(p)$ ”.

A construction pattern can be specified as a normal function invocation like $fun(t)$ or, more often, by embedding certain constant values (such as key strings or values) and necessary notations (such as “.” and “{ }”) into the argument matching term. In the runtime the construction pattern can generate a valid JHM fragment by instantiating the variables with the values they are bound to. For example, the construction pattern “ $people':\{\{ 'surname':\$n \}\}$ ” has the argument term “ $[\$n]$ ”, and it would generate the fragment “ $people':\{\{ 'surname':Li \}; \{ 'surname':Gu \}\}$ ” given the matching result “ $[\$n \rightarrow Li; \$n \rightarrow Gu]$ ”. Besides, function invocations on the subordinate matching terms are allowed to occur in a construction pattern, which makes data construction expressive and compact.

A more common request on data construction is to transform the original extracted data to be of a specified structure rather than follow the original structure in the source document. JPQ can declaratively present the data transformation by employing a deductive restructuring mechanism of matching terms. This mechanism deploys a set of restructuring rules, which constitute a term rewriting system, to indicate the restructuring of matching terms and accordingly direct the transformation of matching results. This method has been successfully utilized in our studies on XML query[13]. In JPQ we substantially modify it to be simple and powerful enough for document database query. The restructuring rules of the matching terms are listed in Fig.1.

These restructuring rules specify the atomic steps of data transformation indicated by restructuring matching terms. For two matching terms t and t' , we

$(t_1, \dots, t_i, t_{i+1}, \dots, t_n) \hookrightarrow (t_1, \dots, t_{i+1}, t_i, \dots, t_n)$	(tuple-commutation)
$(t_1, \dots, t_j, t_{j+1}, \dots, t_n) \hookrightarrow (t_1, \dots, t_j, (t_{j+1}, \dots, t_n))$	(tuple-association)
$(t_1 \dots t_i t_{i+1} \dots t_n) \hookrightarrow (t_1 \dots t_{i+1} t_i \dots t_n)$	(option-commutation)
$(t_1 \dots t_j t_{j+1} \dots t_n) \hookrightarrow (t_1 \dots t_j (t_{j+1} \dots t_n))$	(option-association)
$t \hookrightarrow (t, t)$	(tuple-duplication)
$[t]_i \hookrightarrow \hat{[t]}_i$	(array-flattening)
$(t, t' t'') \hookrightarrow (t, t') (t, t'')$	(option-tuple-distribution)
$(t, [t']_i) \hookrightarrow [(t, t')]_i$ if $\text{var}(t) \cap \text{var}(t') = \emptyset$, $[t']_i$ is not folded array.	(array-tuple-distribution)
$[(t, t')]_i \hookrightarrow [([t]_i, t\%)]_{[t\%]}$	(array-tpl-folding)

Fig. 1. Restructuring rules of matching term

say t is valid (with respect to t'), if t' can be restructured to t by applying the restructuring rules sequentially. For the construction pattern in the construct clause of a JPQ program, it is required that the backbone of the construction pattern is valid with respect to the matching term derived from the extraction pattern. For the valid data transformation indicated by a valid restructured matching term, its restructuring route can be inferred by the deductive mechanism of the rewriting system. Therefore, the requests of data transformation in JPQ can be presented declaratively. Due to space restrictions, we would demonstrate the usage and semantics of the rules with some typical examples. For the formal details and theoretical properties of term restructuring and data transformation, please refer to our technique report[14].

The commutation and the association rules of the tuples and the options are used to rearrange the order of the subordinate terms in the construction pattern. The tuple-duplication rule is used to duplicate data elements to satisfy different construction requests.

As the data elements in a JHM document are often hierarchically organized, it is natural for the extracted data to contain nested arrays. It is a common request to flatten a nested array, that is, to make the elements of the inner array to be the direct elements of the outer array. JPQ introduces a flattened array term $\hat{[t]}_i$ and specifies the array-flattening rule to transform an array $[t]_i$ to a flattened array $\hat{[t]}_i$ if $[t]_i$ is a member of an outer array like $[[t]_i]_{i'}$. After the transformation, the elements in the inner arrays of $[[t]_i]_{i'}$ would be released and directly belong to the outer array.

JPQ allows a tuple containing an option or array to be restructured to an option or an array, as shown in the two distribution rules. This kind of transformation is named distribution because it works like a distribution law in algebra. An option distribution is to restructure the tuple $(t, t_1 | t_2)$ to the option $(t, t_1) | (t, t_2)$ so that the values bound to t can be embedded in different construction pattern depending on its associating pattern t_1 or t_2 . An array distribution is to restructure the tuple $(t, [t']_i)$ to the array $[(t, t')]_i$, that is, to distribute the values of t' to be coupled with the value of t and thus form a new array $[(t, t')]$. In the new array $[(t, t')]$, each element corresponds to an element in the array $[t']_i$ and thus the restructured array can be indexed by the original index term

i. The information of the index array of an array is very helpful for inferring the provenance of array distribution. In a construction pattern an array $[t]_i$ is presented as $[t]$ **groupby** i whereas it is often simplified as $[t]$ if the index term i can be inferred.

Example 1. Find all the presidents and their associating object values, and generate a “presidents” array containing each president’s role and the corresponding object value.

```
from doc("univ")/($r"?president?"):(<$po,{>)|[$pa])
construct {"presidents":[{"role":$r,"info":$po}|^{"role":$r,"info":$pa}]}}
```

As shown in Example 1, the key-value pairs of the presidents are enumerated, the object value would be bound to “\$po” directly and the array value would be further matched with “[\$pa]”. In the construction pattern, the raw term “(\$r,\$po|[\$pa])” is restructured to the term “(\$r,\$po)|(\$r,\$pa)” in which \$r would be used in different ways. Further, the tuple “(\$r,\$pa)” is restructured to the array “[\$r,\$pa]” and then flattened as “^[\$r,\$pa]”. Therefore, in the construction result the presidents’ objects and the vice-presidents’ array members would be merged into a homogeneous array.

Besides indicating the restructuring provenance, an index array can also be used to reorder the elements in the array it refers to. JPQ allows the matching term in an index array to be used as function argument and introduces the “**asc**” and “**desc**” suffix to sort the function values of the elements in the index array. The elements in the value array are accordingly reordered.

Grouping elements by a set of specified index is a common request in data query. JPQ implements this request with folded arrays and the array-tuple-folding rule. The tuple elements in the array $[(t,t')]_i$ can be grouped as an array of equivalence classes on the values of the term t which is denoted as the folded array $[[(t,t')]_i, t\%]$. Naturally, the folded array has the index term $t\%$, indicating the distinct values of the array $[t]$. In the construction patterns, the index pattern $t\%$ is also presented as “**groupby** $t\%$ ” and is often omitted if $t\%$ is used as an indication.

Example 2. For each faculty member list the schools in which he/she has an occupation, order the faculty members in the ascending order of their IDs.

```
from doc("univ"){ "schools":[{"name":$n, "faculty":[{"ID":$id}]}]}
construct {"faculty":[{"ID":^[$id]%, ["school":$n]}] groupby ^[id]% asc}
```

As shown in Example 2, the nested array denoted by “ $[(\$n, \$id)]$ ” is firstly flattened as “ $[(\$n, ^[\$id])]$ ” and then be classified as “ $[[(\$n, ^[\$id])], ^[\$id]\%]$ ” which is used in the construction pattern by hiding the term $^[\$id]$ in the inner arrays. The index term $^[\$id]\%$ has the ordering suffix “**asc**”, indicating to order the array elements by the ascending order of ID.

3.4 Data Filtering

JPQ uses predicate conditions and compound conditions in the *where* clause to filter out the unwanted data elements. Predicate conditions are essentially

predicate function invocations to filter the values bound to the variables in the argument term. A predicate condition can be a simple predicate function invocation, a quantified condition or a composite condition combining subordinate conditions with the boolean connectives “**and**”, “**or**” and “**not**”. The simple predicate function invocation is in the canonical form as $fun(p)$, or is in the infix form like “ $\$a1 = \$a2$ ” for some binary functions.

Example 3. Find the pairs of names of the schools which share common faculty member.

```
from doc("univ"){schools:<[{"name":$n1, "faculty":[{"ID":$id1}]}],
                    [{"name":$n2, "faculty":[{"ID":$id2}]}]>}
construct {"result":[(^ [{"school1":$n1, "school2":$n2}])]}
where not($n1 = $n2) and $id1=$id2
```

In this example, a self-join on the school arrays is established with the composite condition. As the two simple conditions respectively have the arguments “ $(\$n1, \$n2)$ ” and “ $(\$id1, \$id2)$ ”, the composite condition’s argument is the tuple term “ $(\$n1, \$n2, \$id1, \$id2)$ ”.

Quantified conditions are the ones with the quantifiers “*foreach*” or “*for-some*”, in the form of “*quantifier mterm in array; predicate-condition*”, indicating that each (or some) of the element(s) bound to the term in the array should satisfy the predicate condition in the quantified condition body. The quantified term “*p in [p]*” can be simplified as *p*.

Example 4. Find the schools whose faculty amount is larger than 100 or each member has an email address.

```
from doc("univ"){schools:[($s {"faculty":[$f]})]}
construct "result":[$s]
where count[$f]>100 or (foreach $f; notnull($f."email"))
```

A quantified condition’s argument is the array it ranges in rather than the array elements. That is, the two subordinate conditions all have the argument term “ $[f]$ ” and thus the composite condition also has the argument “ $[f]$ ”.

Since different predicate conditions are used to filter the data in different part of the data structure, JPQ adopts compound conditions, the predicate conditions combined with the connectives “**par**” and “**with**”, to form a consistent global requirement of filtering the extracted data. A compound condition c_1 *par* c_2 is a disjunctive condition where c_1 and c_2 are two conditions filtering values of different parts of an option. In this condition the two conditions work in parallel and the filtering results are merged as the final result of the whole condition.

Example 5. For each president find the schools he/she is also a faculty member.

```
from doc("univ")</"?president?":(($p1{"ID":$id1})|(($p2{"ID":$id2}))),
                    {"schools":[{"name":$n, "faculty":[{"ID":$id3}]}]}>
construct {"results":[^ [{"president":$p1, "school":$n}] |
                    ^ [{"president":$p2, "school":$n}]]]}
where id1=id3 par id2=id3
```

In this example, each of the two predicate conditions specifies one branch of the option term “($\$id1, \$id3$)|($\$id2, \$id3$)”. Therefore, we need to combine them with the “par” operator to gather the results filtered by two conditions. A common illusion is to use the operator “or” instead of “par”. However, using “or” means that the composite condition can be treated as a predicate with the argument term “($\$id1, \$id2, \$id3$)”, which is not a valid function in the JPQ language.

A compound condition c_1 with c_2 is used to filter hierarchical data elements. In the runtime, the predicate condition c_2 works on the result data elements filtered by the predicate c_1 . This condition differs from “ c_1 and c_2 ” in that it is processed sequentially rather than commutatively, which indicates a bottom-up order on data hierarchy.

Example 6. Find the schools whose faculty contains at least 100 members who have an email address with an “edu” suffix.

```
from doc("univ"){ "schools": [{ "name": $n, "faculty": [{ "email": $m }] } ] }
construct { "result": [{ "school": $n }] }
where endWith($m, "edu") with count([{$m}]) >= 100
```

In this example, the condition to filter the faculty members and the one to filter the arrays of the faculties work at different levels in the data hierarchy, and thus a bottom-up order to process the data filtering is required.

Data filtering in JPQ is a very subtle issue, because the soundness and the consistency of the semantics for filtering elements in the hierarchical structure and the optional structure are quite complex. For theoretical details of data filtering on these complex structures, please refer to our technique report[14].

4 Conclusion

Document databases are becoming popular in practical data management. However, existing query mechanisms or query languages for JSON-like documents are not expressive enough to present complex document queries in practice. In this paper, we introduced a new query language named JPQ. JPQ is a pattern-based functional language which adopts various expressive patterns to extract structural data elements from JSON-like documents and to construct document fragments based on a deductive mechanism on data transformation. In comparison with the other query languages for JSON-like documents, JPQ exhibits many expressive and interesting features by the coherent pattern-based mechanisms in data extraction, data transformation and data filtering, as Fig.2 shows.

Our study on the JPQ language is still in its preliminary stage. Currently we have finished the design of the core language, and a prototype based on the operational semantics is being implemented. However, although a sound semantics has been developed, how to efficiently process complex queries with compound conditions is still a tough problem to be solved.

The study on some deeper topics of JPQ is also underway. We are extending the language with the update function so as to make it a full-fledged manipulation language. Meanwhile, processing JPQ queries in parallel databases, especially on a Map/Reduce framework, is also an interesting topic that we concern.

Languages					
Features	JPQ	Jaql	JSONiq	UnQL	Query API (MongoDB)
Full-fledged query language?	Yes	Yes	Yes	Yes	No
Data extraction approach	Pattern Matching	Function pipeline	Navigational function	Navigational operator	Function pipeline
Relationships among data elements	Conjunctive, Disjunctive, Aggregation	Not supported	Conjunctive	Conjunctive	Not supported
Schema-less query	Yes	No	Yes	Yes	Yes
Key information query	Pattern matching	Not supported	Special function	Not supported	Not supported
Enumerating value content	Pattern matching	Not supported	Special function	Not supported	Special function
Handling heterogeneous data	Pattern matching	Special statement	Special statement	Not supported	Host program
Hierarchical data construction	Pattern matching	Nested program	Nested program	Not supported	Host program
Deductive data Transformation	Term rewriting system	Not supported	Not supported	Not supported	Not supported
Composite data filtering	Hierarchically pattern filtering	Not supported	Not supported	Not supported	Not supported

Fig. 2. Comparison of JSON-like document data query languages and interfaces

References

1. D. Crockford. The application/json Media Type for JavaScript Object Notation (JSON), RFC 4627, <http://www.ietf.org/rfc>, 2006.
2. N. Nursetov, et-al. Comparison of JSON and XML Data Interchange Formats: A Case Study. *Caine 9* (2009): 157-162.
3. Chodorow, Kristina, and Michael Dirolf. *MongoDB: the definitive guide*. O'Reilly Media, 2010.
4. Anderson, J. Chris, Jan Lehnardt, and Noah Slater. *CouchDB: The Definitive Guide: Time to Relax*. O'Reilly Media, 2010.
5. Orientdb. <http://www.orientdb.org/>
6. RavenDB. <http://http://ravendb.net/>
7. K.Beyer, V. Ercegovic, R. Gemulla, A. Balmin, et-al. Jaql: A scripting language for large scale semistructured data analysis. *Proceedings of the VLDB Endowment* 4, no. 12, 2011.
8. Robie, Jonathan, Matthias Brantner, Daniela Florescu, et-al. XQuery for JSON, JSON for XQuery. *XML Prague 2012* (2012): 63.
9. JPath. <http://bluelinecity.com/software/jpath/>
10. JsonPath. <http://goessner.net/articles/JsonPath/>
11. JsonSelect. <http://jsonselect.org/>
12. Prestegarden, Dag Olav. *UnQL: A Query Language for NoSQL Databases*. PhD diss., Norwegian University of Science and Technology, 2012.
13. X. Li, M. Liu, and Y. Zhang. Towards a "More Declarative" XML Query Language. *Proc. of DEXA'10, LNCS 6262*, pp.375-390, 2010.
14. X. Li, M. Liu. *XTQ: A Declarative Functional XML Query Language*. Technical Report, <http://www.skls.org/xtq>, 2013.